arXiv:1605.06718v1 [stat.ME] 22 May 2016

The de-biased Whittle likelihood for second-order stationary stochastic processes

Adam M. Sykulski*
Department of Statistical Science, University College London. UK
and
Sofia C. Olhede
Department of Statistical Science, University College London, UK
The Alan Turing Institute, UK
and
Jonathan M. Lilly
NorthWest Research Associates, Seattle, USA

May 24, 2016

Abstract

The Whittle likelihood is a computationally efficient pseudo-maximum likelihood inference procedure which is known to produce biased parameter estimates for large classes of time series models. We propose a method for de-biasing Whittle likelihood parameter estimates for second-order stationary stochastic processes. We demonstrate how to compute the de-biased Whittle likelihood in the same $\mathcal{O}(n \log n)$ computational efficiency as standard Whittle likelihood. We prove that the method is consistent, and demonstrate its superior performance in simulation studies. We also demonstrate how the method can be easily combined with standard methods of bias reduction, such as tapering and differencing, to further reduce bias in parameter estimates.

Keywords: Periodogram; Aliasing; Blurring; Tapering; Differencing; Matérn process

*The work of A. M. Sykulski was supported by a Marie Curie International Outgoing Fellowship within the 7th European Community Framework Programme.

1 Introduction

The Whittle likelihood (Whittle 1953) is a pseudo-maximum likelihood estimator, which is commonly used for estimating the parameters of stochastic processes from realized temporal or spatial observations (Pawitan and O'Sullivan 1994, Fuentes 2007, Matsuda and Yajima 2009, Wang and Xia 2015). The method is popular due to its computational efficiency, using Fast Fourier Transforms (FFTs) to approximate the likelihood in $\mathcal{O}(n \log n)$ operations, where n is the length of the observed series. This is in contrast to maximum likelihood which requires the inversion of a covariance matrix, in general a $\mathcal{O}(n^3)$ operation, or $\mathcal{O}(n^2)$ with a regularly sampled stationary process as the covariance matrix is Toeplitz.

The Whittle likelihood is known to commonly produce biased parameter estimates for finite sample sizes (Dahlhaus 1988, Velasco and Robinson 2000, Contreras-Cristan et al. 2006). This occurs because the Whittle likelihood uses the periodogram—a naïve and biased spectral density estimate; and this bias will translate into the parameter estimates. The bias in the periodogram is attributed to *blurring*, sometimes also referred to as *leakage* (Percival and Walden 1993), which occurs because the spectral density estimate is from a finitely-observed sample, where this sample can be viewed as a truncation of an infinite sample. The edge-effects of this truncation cause energy in the spectral density estimate to "leak" from high to low energy regions in the frequency domain. Typically proposed solutions to reducing bias in parameter estimates are to taper or difference the observed series (Dahlhaus 1988, Velasco and Robinson 2000), but these operations come at the expense of reducing the degrees of freedom in the data, and it is not always clear how much one should taper or difference. We propose an additional solution, which is to quantify the expected artifacts from blurring directly into the theoretical spectral density, and then compare this quantity with the observed spectral density estimate in the likelihood.

For continuous stochastic processes, there is the added problem of aliasing (see also Fuentes (2007)), which is attributed to the effect of finitely sampling the observations at fixed time periods. Using the Whittle likelihood to fit such observations to the spectral density of the continuous process creates a further source of bias, as the spectral density estimate will be contaminated by frequencies higher than the observed sampling rate. Our de-biased approach on the other hand naturally accounts for the effects of aliasing and blurring in one operation, while retaining the same $\mathcal{O}(n \log n)$ computational efficiency as standard Whittle likelihood. We note that accounting for aliasing (and not blurring) in the standard Whittle likelihood will in general require a numerical approximation to an integral, to fold in energy from high frequencies into the modeled spectrum. In such cases, the Whittle likelihood will become slower to implement than our de-biased method.

We prove consistency of our method, under more relaxed assumptions than are required for standard Whittle likelihood, as differentiability of the spectrum is now no longer strictly required. This comes at the cost of a reduced rate of convergence, namely $\mathcal{O}(n^{-1/3})$ relative to the normal rate of $\mathcal{O}(n^{-1/2})$. We argue however that the de-biased Whittle is preferred in application due its superior finite-sample performance in terms of bias reduction, and not for its asymptotic properties where Whittle likelihood is already known to be efficient (Dzhaparidze and Yaglom 1983). We demonstrate through simulation studies, using the Matérn process (Gneiting et al. 2010), that the observed bias in parameter estimates for moderate sample sizes is often orders of magnitude lower with the de-biased method, while the variance remains the same, such that the overall estimation error is much reduced.

We demonstrate how our de-biasing method can be efficiently combined with tapering and/or differencing procedures, and show how applying the de-biased form of the Whittle likelihood virtually always reduces the average bias and error of parameter estimates, no matter what choice of taper or differencing operator is made, and often by a large amount.

In this sense, our results build on those of Velasco and Robinson (2000), who advocate the use of tapering and differencing procedures when blurring effects are most pronounced, specifically when the rate of spectral decay is faster than ω^{-2}, where ω denotes frequency.

Small sample effects observed when using the Whittle likelihood have also been explored by Dahlhaus (1988). Dahlhaus showed inconsistency of both standard Whittle estimates, as well as conditional likelihood estimates, when the characteristic root of the time series is approaching unity at the rate $1/n$. Tapered estimates were shown not to exhibit the same effect. Our results correct for bias, the cause of the inconsistency observed in Dahlhaus (1988), whilst outperforming tapered estimators.

Other methods have been proposed for bias correction, see Taniguchi (1983). Taniguchi assumes that the time series is generated by an autocovariance that has the form of a parametric class of models, where the decay of the autocovariance is sufficiently swift. The form of the bias correction requires analytic computation, and is only designed for estimating a scalar parameter. The ideas proposed are interesting, but are not automated and generalized unlike ours, and require user intervention. There have also been alternative pseudo-maximum likelihood methods proposed in the time domain (e.g. Anitescu et al. (2012)), which use circulant embedding to achieve $\mathcal{O}(n \log n)$ computational efficiency, and we contrast with this approach in a simulation study.

The paper is organized as follows. In Section 2 we provide necessary preliminaries for stationary processes. Section 3 formalizes maximum and Whittle likelihood for second-order processes. Section 4 introduces the de-biased Whittle likelihood. We then discuss theoretical properties of our estimator in Section 5. Section 6 incorporates differencing and tapering into the de-biased approach. In Section 7 we perform simulation studies comparing with the state-of-the-art. Concluding remarks are in Section 8. The appendix contains the proof of consistency for the de-biased Whittle likelihood.

2 Preliminaries

2.1 Assumptions

In this paper we assume the stochastic process of interest is modeled in continuous time, however, we present our notation and equations in such a way that results can be easily adjusted for discrete processes. Continuous processes are considered in order that we may address bias effects from both aliasing and blurring. With discrete processes aliasing is no longer an issue, assuming the discrete process is not sub-sampled. Blurring is still very much an issue however, and the de-biased Whittle likelihood can be used to remove bias effects from blurring in this case, in exactly the same way as for continuous processes.

We assume that the Fourier transform of the process has a jointly Gaussian distribution, in order to prove asymptotic consistency of our proposed inference method. It is important to point out that this assumption is not as strict as it may seem. Processes that are non-Gaussian in the time domain may in fact have Fourier transforms with approximately Gaussian distributions for sufficiently large sample size. This is a consequence of a central limit theorem, see Brillinger (2001), who also provides formal conditions when the Gaussian assumption is asymptotically valid. Serroukh and Walden (2000) provide practical examples which satisfy such conditions.

2.2 Definitions and Notation

We define $\{X_t\}$ as the infinite sequence obtained from a zero-mean continuous-time process $X(t)$, that is $X_t \equiv X(t\Delta)$, where $\Delta > 0$ is the sampling interval and $t \in \mathbb{Z}$. Assuming that the process is second-order stationary, we define the autocovariance sequence by $s_X(\tau) = \mathbb{E}\{X_t X_{t+\tau}\}$ for $\tau \in \mathbb{Z}$, where $\mathbb{E}\{\cdot\}$ is the expectation operator. The power spectral density

of $\{X_t\}$ forms a Fourier pair with the autocovariance sequence, and is defined via

$$S_X(\omega) = \Delta \sum_{\tau=-\infty}^{\infty} s_X(\tau) \exp(-\mathrm{i}\omega\tau\Delta), \quad s_X(\tau) = \frac{1}{2\pi} \int_{-\pi/\Delta}^{\pi/\Delta} S_X(\omega) \exp(\mathrm{i}\omega\tau\Delta) d\omega. \tag{2.1}$$

The angular frequency $\omega \in [-\pi/\Delta, \pi/\Delta]$ is given in radians. For a length n sample $\{X_t\}_{t=1}^n$, a simple, but statistically inconsistent, estimate of $S_X(\omega)$ is the periodogram—denoted $\hat{S}_X(\omega)$—which is the squared absolute value of the Discrete Fourier Transform (DFT) defined as

$$\hat{S}_X(\omega) = |J_X(\omega)|^2, \quad J_X(\omega) = \sqrt{\frac{\Delta}{n}} \sum_{t=1}^{n} X_t \exp(\mathrm{i}\omega t\Delta), \quad \omega \in [-\pi/\Delta, \pi/\Delta]. \tag{2.2}$$

Using the Cramér spectral representation theorem we can write X_t in terms of orthogonal increments $\{d\Psi_X(\omega)\}$ where

$$X_t = \int_{-\pi/\Delta}^{\pi/\Delta} d\Psi_X(\omega) \exp(\mathrm{i}\omega t\Delta), \tag{2.3}$$

where $\mathrm{var}\{d\Psi_X(\omega)\} = \frac{1}{2\pi} S_X(\omega) d\omega$. Note that because $\{X_t\}$ is a discrete sequence, $d\Psi_X(\omega)$ has already been aliased. Thus there may be departures between $S_X(\omega)$ and the continuous-time process spectral density, which we denote as $\widetilde{S}_X(\omega)$, and can be defined for $\omega \in \mathbb{R}$ by

$$\widetilde{S}_X(\omega) = \int_{-\infty}^{\infty} \widetilde{s}_X(\lambda) \exp(-\mathrm{i}\omega\lambda) d\lambda, \quad \widetilde{s}_X(\lambda) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \widetilde{S}_X(\omega) \exp(\mathrm{i}\omega\lambda) d\omega, \tag{2.4}$$

where $\widetilde{s}_X(\lambda) = \mathbb{E}\{X(t)X(t-\lambda)\}$ (for $\lambda \in \mathbb{R}$) is the continuous-time process autocovariance, which is related to $s_X(\tau)$ via $\widetilde{s}_X(\tau\Delta) = s_X(\tau)$ when $\tau \in \mathbb{Z}$. It follows that

$$S_X(\omega) \equiv \sum_{k=-\infty}^{\infty} \widetilde{S}_X\left(\omega + k\frac{2\pi}{\Delta}\right). \tag{2.5}$$

for $\omega \in [-\pi/\Delta, \pi/\Delta]$. Thus contributions to $\widetilde{S}_X(\omega)$ outside of the range of frequencies $\pm\pi/\Delta$ are said to be "folded" or "wrapped" into $S_X(\omega)$. We have defined both $S_X(\omega)$ and

$\widetilde{S}_X(\omega)$ here, as both quantities are important in separating the contributions of aliasing and blurring in spectral estimation.

3 Maximum Likelihood and the Whittle Likelihood

Consider a sample $\mathbf{x} = \{X_t\}_{t=1}^n$ observed from a zero-mean Gaussian process, $X(t;\boldsymbol{\theta})$, where $\boldsymbol{\theta}$ is a length-p vector, and $C_X(\boldsymbol{\theta}) \equiv \mathbb{E}\left(XX^T\right)$ is the $n \times n$ theoretical covariance matrix that $\mathbf{x}$ would take, under the assumption that it is a sample drawn from the proposed model. Exact maximum likelihood inference can be performed by evaluating the log-likelihood given by

$$\ell(\boldsymbol{\theta}) = -\frac{1}{2}\log|C_X(\boldsymbol{\theta})| - \frac{1}{2}\mathbf{x}\, C_X^{-1}(\boldsymbol{\theta})\, \mathbf{x}^T, \tag{3.1}$$

where $\mathbf{x}^T$ denotes the transpose of $\mathbf{x}$, and the superscript "-1" is the matrix inverse. The determinant of $C_X(\boldsymbol{\theta})$ is denoted by $|C_X(\boldsymbol{\theta})|$. We have removed additive constants not affected by $\boldsymbol{\theta}$ in (3.1). The optimal choice of $\boldsymbol{\theta}$ for our chosen model to characterize the observed data is then found by maximizing the likelihood function (Brockwell and Davis 2009)

$$\widehat{\boldsymbol{\theta}} = \arg\max_{\boldsymbol{\theta}\in\boldsymbol{\Theta}} \ell(\boldsymbol{\theta}).$$

Because the time-domain maximum likelihood is known to have optimal properties, any other estimator will be compared with the properties of this quantity.

A standard technique to avoiding expensive matrix inversions, is to approximating equation (3.1) in the frequency domain, following the seminal work of Whittle (1953). This approach approximates $C_X(\boldsymbol{\theta})$ using a Fourier representation, and utilizes the special properties of Toeplitz matrices. For a single series of observations the Whittle likelihood,

denoted $\ell_W(\boldsymbol{\theta})$, once discretized, is given by

$$\ell_W(\boldsymbol{\theta}) = -\frac{1}{2}\sum_{\omega\in\Omega}\left[\log\left\{\widetilde{S}_X(\omega;\boldsymbol{\theta})\right\} + \frac{\hat{S}_X(\omega)}{\widetilde{S}_X(\omega;\boldsymbol{\theta})}\right], \tag{3.2}$$

where $\widetilde{S}_X(\omega;\boldsymbol{\theta})$ is the parametric form of the theoretical spectrum of the continuous-time process, defined in equation (2.4), and Ω is the set of discrete Fourier frequencies: $\frac{2\pi}{n\Delta}(-\lceil\frac{n}{2}\rceil+1,\ldots,-1,0,1,\ldots,\lfloor\frac{n}{2}\rfloor)$. The subscript "$W$" in $\ell_W(\boldsymbol{\theta})$ is used to denote "Whittle." We note that sometimes this summation is made over positive or nonnegative frequencies only, which then drops the factor of $1/2$ in equation (3.2). Asymptotically this makes no difference, however for finite samples, it is better to sum over both positive and negative Fourier frequencies, thus ensuring the degrees of freedom in the periodogram are correctly aggregated.

The Whittle likelihood *approximates* the time-domain likelihood when $\Omega \subset \left[-\frac{\pi}{\Delta},\frac{\pi}{\Delta}\right]$, i.e. $\ell(\boldsymbol{\theta}) \approx \ell_W(\boldsymbol{\theta})$, and this statement can be made precise, see Dzhaparidze and Yaglom (1983). However its computational cost is $O(n\log n)$ versus $O(n^2)$ for the time-domain likelihood. We note there have also been recent advances in constructing $O(n\log n)$ time-domain approximations to maximum likelihood (see Anitescu et al. (2012)), and we compare performance in a simulation study in Section 7.

4 The De-Biased Whittle Likelihood

The Whittle likelihood utilizes the periodogram, $\hat{S}_X(\omega)$, which is an inconsistent and biased measure of the continuous time process' spectral density, due to the blurring caused by blurring and aliasing effects (Percival and Walden 1993). Aliasing results from the discrete sampling of the continuous-time process to generate an infinite sequence, whereas blurring is associated with the truncation of the infinite sequence $\{X_t\}$ over a finite time interval.

The desirable properties of the Whittle likelihood rely on the *asymptotic* behavior of the periodogram for very large sample sizes. The bias of the periodogram for finite samples however, will translate into biased parameter estimates of the Whittle likelihood, as has been widely reported (see e.g. Dahlhaus (1988)). We therefore define an alternative pseudo-maximum likelihood function given by

$$\ell_D(\boldsymbol{\theta}) = -\frac{1}{2}\sum_{\omega\in\Omega}\left[\log\left\{\overline{S}_X(\omega;\boldsymbol{\theta})\right\} + \frac{\hat{S}_X(\omega)}{\overline{S}_X(\omega;\boldsymbol{\theta})}\right], \tag{4.1}$$

$$\overline{S}_X(\omega;\boldsymbol{\theta}) = \int_{-\pi/\Delta}^{\pi/\Delta} S_X(\nu;\boldsymbol{\theta})\mathcal{F}_{n,\Delta}\left(\nu-\omega\right)\, d\nu, \quad \mathcal{F}_{n,\Delta}(\omega) = \frac{\Delta}{2\pi n}\frac{\sin^2(n\omega\Delta/2)}{\sin^2(\omega\Delta/2)}, \tag{4.2}$$

where the subscript "D" stands for "de-biased." Here $\overline{S}_X(\omega;\boldsymbol{\theta})$ is the *expected* periodogram, which is the convolution of the true modeled spectrum with the Fejér kernel $\mathcal{F}_{n,\Delta}(\omega)$, such that $\mathbb{E}\{\hat{S}_X(\omega)\} = \overline{S}_X(\omega;\boldsymbol{\theta})$ (Bloomfield 1976). We call $\overline{S}_X(\omega;\boldsymbol{\theta})$ the *blurred* spectrum and equation (4.1) the *de-biased* Whittle likelihood. The set Ω is defined as in equation (3.2).

While the concept of using the blurred spectrum $\overline{S}_X(\omega;\boldsymbol{\theta})$ in equation (4.1) is simple, the key innovation of our method lies in how it is efficiently computed without losing the $\mathcal{O}(n\log n)$ computational efficiency of the likelihood estimator. If we directly use equation (4.2), then this integral would usually need to be approximated by discretization, and could be computationally expensive. However we employ the useful trick that a frequency domain convolution can be converted exactly into a time domain multiplication. Specifically, $\overline{S}_X(\omega;\boldsymbol{\theta})$ can be efficiently computed by working with the Fourier pair of the blurred spectrum, which is the expectation of the *biased autocovariance estimator*, defined by $\overline{s}_X(\tau;\boldsymbol{\theta}) = \mathbb{E}\{\frac{1}{n}\sum_{t=1}^{n-|\tau|} X_t X_{t+\tau}\}$ for $\tau = 0, \pm 1, \pm 2 \ldots, \pm(n-1)$, such that

$$\overline{S}_X(\omega;\boldsymbol{\theta}) = \Delta \sum_{\tau=-(n-1)}^{n-1} \overline{s}_X(\tau;\boldsymbol{\theta})\exp(-\mathrm{i}\omega\tau\Delta). \tag{4.3}$$

The sequence $\overline{s}_X(\tau;\boldsymbol{\theta})$ can then be computed exactly from the theoretical autocovariance

sequence in $\mathcal{O}(n)$ operations by

$$\overline{s}_X(\tau;\boldsymbol{\theta}) = \mathbb{E}\left(\frac{1}{n}\sum_{t=1}^{n-|\tau|} X_t X_{t+\tau}\right) = \left(1-\frac{|\tau|}{n}\right) s_X(\tau;\boldsymbol{\theta}). \tag{4.4}$$

Combining equations (4.3) and (4.4), it follows that the blurred spectrum $\overline{S}_X(\omega;\boldsymbol{\theta})$ is exactly given by

$$\overline{S}_X(\omega;\boldsymbol{\theta}) = 2\Delta \cdot \Re\left\{\sum_{\tau=0}^{n-1}\left(1-\frac{\tau}{n}\right) s_X(\tau;\boldsymbol{\theta})\exp(-\mathrm{i}\omega\tau\Delta)\right\} - \Delta \cdot s_X(0;\boldsymbol{\theta}). \tag{4.5}$$

see also (Percival and Walden 1993, p.198). Then to compute the de-biased Whittle likelihood, we only need to evaluate $\overline{S}_X(\omega;\boldsymbol{\theta})$ at the n discrete Fourier frequencies of Ω used in equation (4.1). Therefore $\overline{S}_X(\omega;\boldsymbol{\theta})$ can be exactly computed from $s_X(\tau;\boldsymbol{\theta})$ for $\tau = 0,\ldots,n-1$ using a discrete Fourier transform in $O(n\log n)$ operations, where care must be taken to subtract the variance term, $\Delta \cdot s_X(0;\boldsymbol{\theta})$, to avoid double counting the main diagonal in the covariance matrix. This efficient computation been made possible by transforming the frequency-domain convolution in equation (4.2) into a time-domain multiplication, where the multiplication involves combining the autocovariance sequence with the triangle kernel $\left(1-\frac{\tau}{n}\right)$. Both aliasing and blurring effects are automatically accounted for in equation (4.5); the effect of *aliasing* is accounted for by sampling the theoretical autocovariance function at discrete times, while the effect of *blurring*, due to the truncation of the sample to finite length, is accounted for by the triangle kernel.

4.1 Discussion

One of the useful features of the de-biased Whittle likelihood is that it can be directly computed from the model for the autocovariance of the stochastic process, $s_X(\tau)$, in $\mathcal{O}(n\log n)$ operations, without having to derive the analytical form of the spectral density, $\widetilde{S}_X(\omega)$.

This makes the method are more directly applicable alternative to the standard maximum likelihood approach of (3.1). If the closed form of the autocovariance sequence is unknown however, then we can inverse Fourier transform the theoretical spectrum of the process to approximate the autocovariance sequence numerically (by discretizing the integral in equation (2.4)), where this approximation can be made more accurate by oversampling the spectrum. Subsequently we would compute the discrete Fourier transform of this sequence, evaluated at the observed sample lags, multiplied by the triangle kernel (as in equation (4.5)), to recover an estimate of $\overline{S}_X(\omega)$, still in $\mathcal{O}(n \log n)$ operations.

Computing the Whittle likelihood with the aliased, but not blurred, spectrum $S_X(\omega; \boldsymbol{\theta})$—defined in equation (2.1)—is more complicated, as this seldom has an analytic form for continuous processes, and must be instead approximated by either explicitly wrapping in contributions from frequencies higher than the Nyquist as in (2.5), or via an approximation to the Fourier transform in equation (2.1). This is in contrast to the de-biased Whittle likelihood, where the effect of aliasing and blurring can be computed exactly in one operation, as in equation (4.5). The de-biased Whittle likelihood is therefore generally quicker to implement than standard Whittle likelihood using an aliased spectrum.

5 Properties of the De-Biased Whittle Likelihood

It is known that standard Whittle likelihood provides a consistent estimator of the spectrum (Dzhaparidze and Yaglom 1983). In this section, we establish consistency and related properties of the de-biased Whittle likelihood. The main difficulty in the proof is that, although the de-biased Whittle likelihood accounts for the bias of the periodogram, there is still present the broadband correlation between frequencies of the periodogram caused by the Fejér kernel. This is what prevents the de-biased Whittle likelihood being exactly

equal to the time-domain maximum likelihood for Gaussian data.

To establish consistency, we need to bound the asymptotic behavior of this correlation. The statement is provided below in Theorem 1, with the proof provided in the appendix. The proof is composed of three propositions which provide bounds, in turn, for the covariance of the DFT, the variance of linear combinations of the periodogram, and finally the score and Hessian of the de-biased Whittle likelihood. Together these establish that the de-biased Whittle likelihood is a consistent estimator converging with probability, at a rate $n^{-1/3}$. This theorem requires weaker assumptions than for standard Whittle likelihood, as it does *not* require the assumptions that the spectrum is differentiable in ω and is near-constant over the width of the Fejér kernel.

Theorem 1. *Assume that the infinite sequence $\{X_t\}$ is a zero-mean second-order stationary discrete process where $\overline{S}_X(\omega;\boldsymbol{\theta})$, defined in equation (4.1), is twice differentiable in $\boldsymbol{\theta}$. Assume that the spectral density of $\{X_t\}$ is bounded above by the finite value $\|S_X\|_\infty$ and below by the non-zero value $S_{\min}$. Then the estimator*

$$\hat{\boldsymbol{\theta}} = \operatorname{argmax} \ell_D(\boldsymbol{\theta}),$$

for a sample $\{X_t\}_{t=1}^n$, with $\ell_D(\boldsymbol{\theta})$ being the de-biased Whittle likelihood defined in equation (4.1), satisfies

$$\hat{\boldsymbol{\theta}} = \boldsymbol{\theta} + \mathcal{O}_P\left(n^{-1/3}\right).$$

This result requires that the spectrum, $S_X(\omega)$, is bounded from above and below, and that the de-biased Whittle likelihood, $\ell_D(\boldsymbol{\theta})$, is twice differentiable in $\boldsymbol{\theta}$. Standard theory shows that the standard Whittle likelihood approach is consistent with a $n^{-1/2}$ rate if the spectrum is twice differentiable in ω and bounded below and above by a non-zero constant (see Dzhaparidze and Yaglom (1983)). Therefore while the rate of convergence we prove is

slower than that for the standard Whittle likelihood, our method of proof requires weaker assumptions. Furthermore, the key innovation of the de-biased Whittle likelihood is that it performs significantly better than standard Whittle likelihood for finite sample sizes, as we shall demonstrate. In practice, we can see that the variances of the standard and de-biased Whittle estimators will behave similarly, as sums of weighted periodograms will behave similarly whether weighted by quantities involving $\bar{S}_X(\omega)$ or $S_X(\omega)$. We investigate this in more detail in Section 7 through simulation studies.

Theorem 1 establishes that $\hat{\boldsymbol{\theta}}$ is a consistent estimator. Having established consistency, we can use this result to obtain variance estimates of the de-biased Whittle estimators. Note that for simplicity, and without loss of generality, we set $\Delta = 1$ for the remainder of this section. From equations (A.19), (A.20) and (A.21) in the appendix we see that

$$\operatorname{Var}\left\{\hat{\boldsymbol{\theta}}\right\} \sim \boldsymbol{H}^{-1}(\boldsymbol{\theta}) \operatorname{Var}\left\{\nabla \ell_D(\boldsymbol{\theta})\right\} \boldsymbol{H}^{-1}(\boldsymbol{\theta}), \tag{5.1}$$

where $\nabla = \left(\frac{\partial}{\partial\theta_1}, \ldots, \frac{\partial}{\partial\theta_p}\right)$. The matrix $\boldsymbol{H}(\boldsymbol{\theta})$ is defined entrywise by

$$H_{ij}(\boldsymbol{\theta}) = \mathbb{E}\left\{\frac{\partial^2}{\partial\theta_i \partial\theta_j} \ell_D(\boldsymbol{\theta})\right\}$$

and can be approximated—numerically, if the analytic form is not known—by evaluating the Hessian at $\hat{\boldsymbol{\theta}}$. The remaining term, $\nabla\ell_D(\boldsymbol{\theta})$, is the likelihood score, and to estimate its variance we use equation (4.1), expressing the partial derivative in each element of $\boldsymbol{\theta}$ separately to obtain

$$\begin{aligned}\operatorname{Var}\left\{\frac{\partial}{\partial\theta_i}\ell_D(\boldsymbol{\theta})\right\} &= \operatorname{Var}\left\{\sum_{\omega} \frac{\partial \bar{S}_X(\omega;\boldsymbol{\theta})}{\partial\theta_i} \cdot \frac{\hat{S}_X(\omega)}{\bar{S}_X^2(\omega;\boldsymbol{\theta})}\right\} = \operatorname{Var}\left\{\sum_{j=0}^{\lfloor n/2\rfloor} a_{ij}(\boldsymbol{\theta}) \hat{S}_X\left(\frac{2\pi j}{n}\right)\right\} \\ &= \sum_{j=0}^{\lfloor n/2\rfloor}\sum_{k=0}^{\lfloor n/2\rfloor} a_{ij}(\boldsymbol{\theta}) a_{ik}(\boldsymbol{\theta}) \operatorname{Cov}\left\{\hat{S}_X\left(\frac{2\pi j}{n}\right), \hat{S}_X\left(\frac{2\pi k}{n}\right)\right\},\end{aligned}$$

where we have defined

$$a_{ij}(\boldsymbol{\theta}) = \frac{\partial \bar{S}_X\left(\frac{2\pi j}{n};\boldsymbol{\theta}\right)}{\partial \theta_i} \cdot \frac{1}{\bar{S}_X^2(\frac{2\pi j}{n};\boldsymbol{\theta})}.$$

In deriving this expression we have made use of the fact that the $\frac{\partial}{\partial\theta}\log\{\overline{S}_X(\omega;\boldsymbol{\theta})\}$ term is deterministic and therefore does not contribute to the variance. As we have established consistency for $\hat{\boldsymbol{\theta}}$ we can now use the invariance principle of maximum likelihood to construct an estimator of the variance, that is

$$\widehat{\mathrm{Var}}\left\{\frac{\partial}{\partial\theta_i}\ell_D(\boldsymbol{\theta})\right\} = \sum_{j=1}^{\lfloor n/2\rfloor}\sum_{k=1}^{\lfloor n/2\rfloor} \hat{a}_{ij}(\boldsymbol{\theta})\hat{a}_{ik}(\boldsymbol{\theta})\widehat{\mathrm{Cov}}\left\{\hat{S}_X\left(\frac{2\pi j}{n}\right), \hat{S}_X\left(\frac{2\pi k}{n}\right)\right\}.$$

Because $\hat{\boldsymbol{\theta}}$ is consistent we may first take $\hat{a}_{ij}(\boldsymbol{\theta}) = a_{ij}(\hat{\boldsymbol{\theta}})$, and then to estimate the covariance of the periodogram we approximate the integral

$$\widehat{\mathrm{Cov}}\left\{\hat{S}_X(\omega_1), \hat{S}_X(\omega_2)\right\} = \left|\frac{1}{2\pi n}\int_{-\pi}^{\pi} S_X\left(\omega';\hat{\boldsymbol{\theta}}\right) D_n(\omega_1-\omega') D_n^*(\omega_2-\omega')\ d\omega'\right|^2,$$

where $D_n(\omega)$ is the Dirichlet kernel defined by

$$D_n(\omega) = \frac{\sin(n\omega/2)}{\sin(\omega/2)}\exp(-\mathrm{i}\omega(n+1)/2).$$

The off-diagonal terms of $\mathrm{Var}\{\nabla\ell_D(\boldsymbol{\theta})\}$, namely $\mathrm{Cov}\left\{\frac{\partial}{\partial\theta_i}\ell_D(\boldsymbol{\theta}), \frac{\partial}{\partial\theta_j}\ell_D(\boldsymbol{\theta})\right\}$, can be found in exactly the same way. Then substituting into (5.1), along with the Hessian, provides estimates of the variance of the estimators. Normality of $\hat{\boldsymbol{\theta}}$ follows because Proposition 3 in Section S.2 of the appendix shows that $-(1/n)\frac{\partial^2}{\partial\theta_i^2}\ell_D(\boldsymbol{\theta})$ converges in probability to a positive constant, while $(1/n)\frac{\partial}{\partial\theta_i\theta_j}\ell_D(\boldsymbol{\theta})$ is a Gaussian quadratic form where the diagonalized representation does not put too much mass at any individual variate; thus the sum will be a Gaussian random variable.

6 Tapering and Differencing

Tapering and differencing are often suggested as two methods for improving Whittle estimates (Dahlhaus 1988, Velasco and Robinson 2000). Here in this section we outline how the de-biased Whittle likelihood can be combined with either of these procedures.

6.1 Tapered Whittle Inference

To ameliorate spectral blurring of the periodogram, a standard approach is to pre-multiply the data sequence with some weighting function known as a data taper (Thomson 1982). The taper is chosen to have spectral properties such that blurring will be minimized, and the variance of the spectral estimate at each frequency is reduced; however this comes at the expense of increasing correlation between neighboring frequencies.

The tapered Whittle likelihood corresponds to replacing the direct spectral estimator formed from $\hat{S}_X(\omega)$ in equation (2.2) with one using the taper $\boldsymbol{h} = \{h_t\}$

$$J_X(\omega; \boldsymbol{h}) = \sqrt{\Delta} \sum_{t=1}^{n} h_t X_t \exp(-\mathrm{i}\omega t \Delta), \quad \hat{S}_X(\omega; \boldsymbol{h}) = \left|J_X(\omega; \boldsymbol{h})\right|^2, \quad \sum_{t=1}^{n} |h_t|^2 = 1,$$

where h_t is real-valued. Setting $h_t = 1/\sqrt{n}$ recovers the periodogram estimate of equation (3.2). In the parametric setting we then compute

$$\ell_T(\boldsymbol{\theta}) = -\frac{1}{2} \sum_{\omega \in \Omega} \left[\log \left\{ \widetilde{S}_X(\omega; \boldsymbol{\theta}) \right\} + \frac{\hat{S}_X(\omega; \boldsymbol{h})}{\widetilde{S}_X(\omega; \boldsymbol{\theta})} \right]. \tag{6.1}$$

We call this the *tapered Whittle likelihood*, where the subscript "T" denotes that a taper has been used. Velasco and Robinson (2000) demonstrated that for certain discrete processes it is beneficial to use this estimator, rather than the standard Whittle likelihood, for parameter estimation—particular when the spectrum exhibits a large dynamic range.

Nevertheless, tapering in itself will *not* remove all blurring effects, and the issue of aliasing for continuous sampled processes remains.

A useful feature of our de-biasing procedure is that we can combine the method with tapering, to further reduce bias estimates in the maximum likelihood estimation. To do this we define the likelihood given by

$$\begin{aligned} \ell_{TD}(\boldsymbol{\theta}) &= -\frac{1}{2}\sum_{\omega\in\Omega}\left[\log\left\{\overline{S}_X(\omega;\boldsymbol{h},\boldsymbol{\theta})\right\} + \frac{\hat{S}_X(\omega;\boldsymbol{h})}{\overline{S}_X(\omega;\boldsymbol{h},\boldsymbol{\theta})}\right], \qquad (6.2) \\ \overline{S}_X(\omega;h,\boldsymbol{\theta}) &= \int_{-\pi/\Delta}^{\pi/\Delta} S_X(\nu;\boldsymbol{\theta})\mathcal{H}_\Delta\left(\nu-\omega\right)\,d\nu, \quad \mathcal{H}_\Delta(\omega) = \Delta\left|\sum_{t=1}^{n} h_t\exp(-\mathrm{i}\omega t\Delta)\right|^2. \end{aligned}$$

We call this the *tapered de-biased Whittle likelihood*, and it can be computed exactly and efficiently using a similar $\mathcal{O}(n\log n)$ calculation to equation (4.5) to find $\overline{S}_X(\omega;\boldsymbol{h},\boldsymbol{\theta})$ from the time domain

$$\overline{S}_X(\omega;\boldsymbol{h},\boldsymbol{\theta}) = 2\Delta\cdot\Re\left\{\sum_{\tau=0}^{n-1} s_X(\tau;\boldsymbol{\theta})\left(\sum_{t=1}^{n-\tau} h_t h_{t+\tau}\right)\exp(-\mathrm{i}\omega\tau\Delta)\right\} - \Delta\cdot s_X(0;\boldsymbol{\theta}).$$

In equation (6.2), $\hat{S}_X(\omega;\boldsymbol{h})$ is now a tapered spectral estimate, and $\overline{S}_X(\omega;\boldsymbol{h},\boldsymbol{\theta})$ is the theoretical form for the *expected* tapered spectral estimate given $\boldsymbol{\theta}$ and $\boldsymbol{h}$. Accounting for the exact taper used in $\overline{S}_X(\omega;\boldsymbol{h},\boldsymbol{\theta})$ accomplishes debiasing.

We note that the time-domain kernel $\sum_{t=1}^{n-\tau} h_t h_{t+\tau}$ can be pre-computed (unless it has a known analytic form) which requires $\mathcal{O}(n^2)$ operations, but this can be stored for each taper and data-length n, and does not have to be recomputed each time when performing numerical optimization, as it will remain fixed. The computation then involves a Fourier transform after multiplying the taper kernel with the autocovariance sequence. Thus the tapered de-biased Whittle likelihood is still in practice an $\mathcal{O}(n\log n)$ pseudo-maximum likelihood estimator.

With the modifications to the Whittle likelihood we have proposed, the practitioner is free to select between a tapered or periodogram spectral estimate in the same way as before, but can now account for finite sample bias effects by using the de-biased likelihood approximation. Both the de-biased tapered and de-biased periodogram likelihoods have their merits, but these trade-offs are different with *nonparametric* spectral density estimation than they are with *parametric* model estimation. For example, although tapering *decreases* the variance of nonparametric estimates at each frequency, it conversely can *increase* the variance of estimated parameters. This is because the taper is reducing degrees of freedom in the data, which increases correlations between local frequencies. On the other hand, the periodogram creates broadband correlations between frequencies, especially for processes with a high dynamic range. In such instances, even though the de-biased Whittle likelihood accounts for the expected blurring, the broadband correlation can lead to increased parameter errors as compared with tapered (and de-biased) estimates. We explore these effects in greater detail through Monte Carlo simulations in Section 7, and demonstrate that whether one tapers or not, the de-biasing step significantly reduces bias and error.

6.2 Differencing

Another method of reducing the effects of blurring on Whittle estimates is to fit parameters to the differenced process instead. This was illustrated in numerical simulations performed in Velasco and Robinson (2000), where the Whittle likelihood was found to perform poorly with fractionally differenced processes that were more smooth, but improved when working with the differenced process.

Whittle likelihood using the differenced process proceeds as follows. Define $Y_t = X_{t+1} - X_t$, both in terms of the theoretical process and the observed sample. The Whittle

likelihood is then performed by maximizing

$$\ell_W(\boldsymbol{\theta}) = -\frac{1}{2}\sum_{\omega\in\Omega}\left[\log\left\{\widetilde{S}_Y(\omega;\boldsymbol{\theta})\right\} + \frac{\hat{S}_Y(\omega)}{\widetilde{S}_Y(\omega;\boldsymbol{\theta})}\right], \quad \text{where} \quad \widetilde{S}_Y(\omega;\boldsymbol{\theta}) = 4\sin^2\left(\frac{\omega}{2}\right)\widetilde{S}_X(\omega;\boldsymbol{\theta}), \tag{6.3}$$

and $\hat{S}_Y(\omega)$ is the periodogram of the sample, $\{Y_t\}_{t=1}^{n-1}$, where one degree of freedom has been lost by differencing, such that the Fourier frequencies are now $\frac{2\pi}{(n-1)\Delta}(-\lceil\frac{n-1}{2}\rceil + 1, \ldots, -1, 0, 1, \ldots, \lfloor\frac{n-1}{2}\rfloor)$ The de-biased Whittle likelihood is also straightforward to compute from $\{Y_t\}_{t=1}^{n-1}$

$$\ell_D(\boldsymbol{\theta}) = -\frac{1}{2}\sum_{\omega\in\Omega}\left[\log\left\{\overline{S}_Y(\omega;\boldsymbol{\theta})\right\} + \frac{\hat{S}_Y(\omega)}{\overline{S}_Y(\omega;\boldsymbol{\theta})}\right], \tag{6.4}$$

$$\overline{S}_Y(\omega;\boldsymbol{\theta}) = 2\Delta\cdot\Re\left\{\sum_{\tau=0}^{n-1}\left(1-\frac{\tau}{n}\right)s_Y(\tau;\boldsymbol{\theta})\exp(-\mathrm{i}\omega\tau\Delta)\right\} - \Delta\cdot s_Y(0;\boldsymbol{\theta}),$$

where

$$s_Y(\tau;\boldsymbol{\theta}) = 2s_X(\tau;\boldsymbol{\theta}) - s_X(\tau+1;\boldsymbol{\theta}) - s_X(\tau-1;\boldsymbol{\theta}),$$

from direct calculation. This likelihood is still an $\mathcal{O}(n\log n)$ operation to evaluate, as computing $s_Y(\tau;\boldsymbol{\theta})$ is $\mathcal{O}(n)$, and the rest of the calculation is the same as in equation (4.5). Differencing and tapering can also be easily combined, with both the standard and de-biased Whittle likelihood. Furthermore, differencing can be applied multiple times if desired.

To see theoretically how differencing can reduce the variance of the estimators, we explore how the score of the likelihood behaves in the de-biased Whittle likelihood. The score is zero mean (for any finite n), and the variance (as derived in equation (A.11) of the appendix) can be bounded by

$$\mathrm{Var}\left\{\frac{1}{n}\frac{\partial}{\partial\theta_i}\ell_D(\boldsymbol{\theta})\right\} \leq \frac{\|S_X\|_\infty^2\|\frac{\partial\bar{S}_X}{\partial\theta_i}\|_\infty^2}{\bar{S}_{\min}^4}\frac{3}{n^{2/3}}\left\{1+o(1)\right\}, \tag{6.5}$$

where $\|S_X\|_\infty^2$ and $\|\frac{\partial \bar{S}_X}{\partial \theta_i}\|_\infty^2$ are upper bounds on the spectrum and the partial derivative of the blurred spectrum respectively, and $\bar{S}_{\min}$ is a lower bound on the blurred spectrum. The significance of equation (6.5) is that the magnitude of the variance of the score is controlled by the dynamic range of the spectrum. A high dynamic range increases the value of the first ratio in the bound $\|S_X\|_\infty^2 \|\frac{\partial \bar{S}_X}{\partial \theta_i}\|_\infty^2 / \bar{S}_{\min}^4$. This suggests differencing a sampled process with steep spectral slopes, as reducing this ratio will reduce the variance of the estimators, taking care to omit the zero frequency from the fit. Differencing multiple times however will send $\bar{S}_{\min}$ to zero, and at some point the ratio will increase and lead to inferior estimates. We explore the merits of differencing in more detail in the next section.

7 Monte-Carlo Simulations

7.1 Comparing the standard and de-biased Whittle likelihood

In this section we investigate the effectiveness of de-biasing the Whittle likelihood in a Monte Carlo study using data from a Matérn process, comparing across different frequency domain estimators. All MATLAB code to exactly replicate the simulations in this section can be found at `www.ucl.ac.uk/statistics/research/spg/software`.

The Matérn process (Gneiting et al. 2010), is a three-parameter continuous Gaussian stochastic process defined by its spectral density

$$\widetilde{S}_X(\omega) = \frac{A^2}{(\omega^2 + c^2)^\alpha}. \tag{7.1}$$

The parameter A controls the magnitude of the variability, $c > 0$ controls the damping timescale, and $\alpha > 1/2$ controls the rate of spectral decay or equivalently the smoothness or differentiability of the process. When $\alpha > 1$, the power spectrum of the process

will exhibit a high dynamic range, and we can expect the periodogram to be a poor estimator of the spectral density due to blurring. Conversely, when $\alpha < 1$ then we can expect departures between the periodogram and the continuous spectral density because of aliasing. For this reason we will investigate the performance of estimators over a range of α values. We choose to investigate the Matérn as it is a simple yet flexible continuous stochastic process. The Matérn is in some sense a continuous-time analogue of an Autoregressive Fractionally Integrated Moving Average (ARFIMA) model (Granger and Joyeux 1980), where α in the Matérn behaves similarly to the difference parameter (usually denoted as d) in an ARFIMA. This allows us to draw parallels with the simulation findings of Velasco and Robinson (2000), who report large errors in standard Whittle likelihood approximations when $d > 1$.

In Figure 1 we display the bias and standard deviation of the different Whittle estimators for the three parameters $\{A, \alpha, c\}$ where α varies from [0.6,2.5] in intervals of 0.1. We fix $A = 1$ and $c = 0.2$. For each value of α, we simulate 10,000 series of length $n = 1000$, and use these Monte Carlo replicates to calculate bias and standard deviation for each estimator. We implement several different Whittle estimators: standard Whittle likelihood (3.2), tapered Whittle likelihood (6.1), and differenced Whittle likelihood (6.3). In addition, for each of these we implement the de-biased version (equations (4.1), (6.2), and (6.4), respectively). The choice of data taper is the Discrete Prolate Spheroidal Sequence (DPSS) taper (Slepian and Pollak 1961), with bandwidth parameter equal to 4. We note that the performance of tapered versions of the likelihood, relative to other estimators, was found to be broadly similar across different choices of bandwidth. We also performed a combined differenced and tapered version of the standard and de-biased Whittle likelihood, as discussed in Section 6.2, but these results are not included here as performance was virtually identical to tapering without differencing.

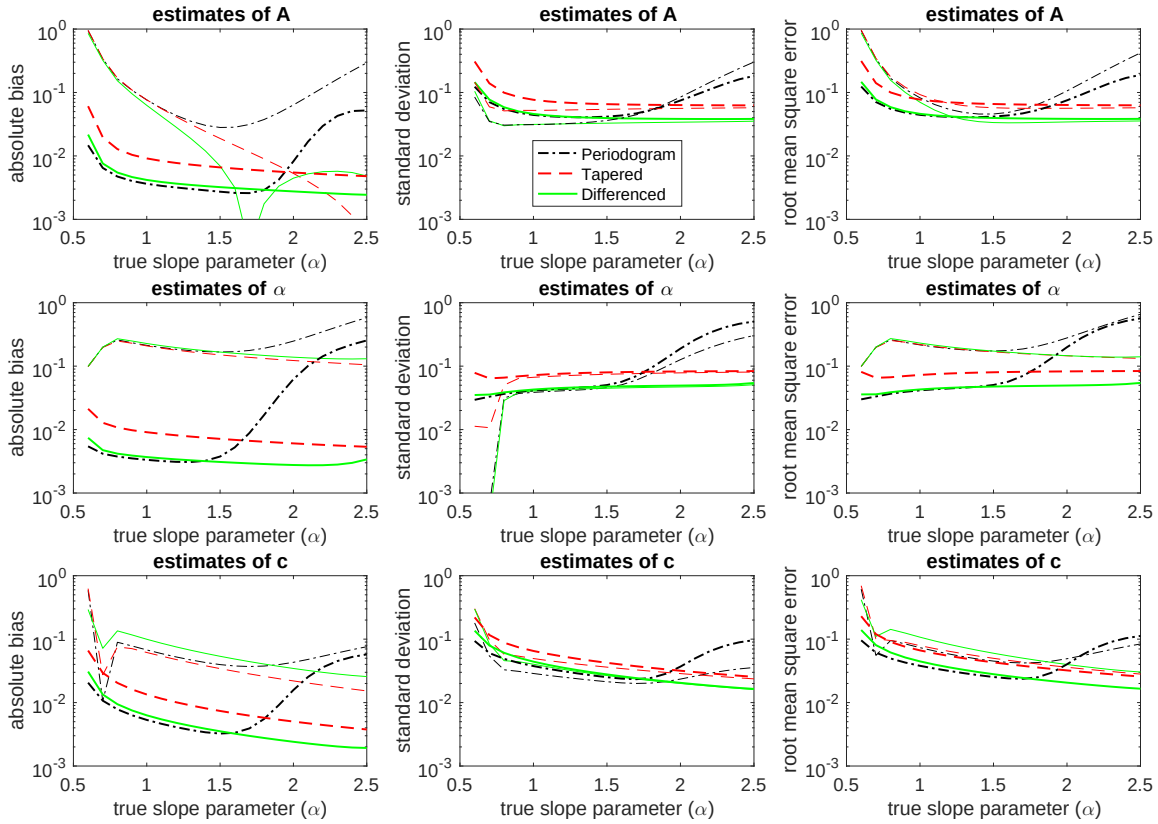

Figure 1: Absolute bias (left column), standard deviation (center column) and root mean square error (right column) of parameter estimates using different forms of the Whittle and de-biased likelihood. The line-styles and colors correspond to different approaches: periodogram, tapered, or differenced (as indicated by the legend in the top center panel). For each color/line-style, the thinner lines are for standard Whittle approaches, and the thicker lines are the corresponding de-biased approaches. The top row corresponds to estimates for the amplitude parameter A, the second row for the slope parameter α, and the bottom row for the damping parameter c as given in equation (7.1). In all panels the results are reported for a range of α values in increments of 0.1 from 0.6 to 2.5, which translates to a spectral slope decaying from rate $\omega^{-1.2}$ to ω^{-5}.

The optimization is performed in MATLAB using `fminsearch`, and uses identical settings for all likelihoods, where initialized guesses for the slope and amplitude are found by performing least squares on the spectral slope, and the initial guess for the damping parameter c is set at a mid-range value of 100 times the Rayleigh frequency (i.e. $c = 100\pi/n$.)

The first column in Figure 1, displays the absolute bias of each type of Whittle estimator for each Matérn parameter. Solid lines represent standard approaches (black for periodogram, red for tapered, and green for differenced); then with dashed-lines we display the performance of the respective de-biased versions. In general using the de-biased method reduces the observed absolute bias with each parameter. Exceptions only occur in two instances where the bias of the tapered and differenced Whittle likelihoods crosses zero when estimating the amplitude parameter A, which causes a "dip" when the absolute value is taken. Note that the absolute biases are displayed on a log10 scale, such that we can in many instances see bias reduction of over a factor of 10, representing over a 90% reduction in bias. The "U" shape over α that we observe with the standard Whittle likelihood using the periodogram, and less so with the tapered method, corresponds to the contamination of aliasing for small α and blurring for large α. Differencing ameliorates the blurring effects for high α, but not the aliasing effects for low α. De-biased methods, particularly when combined with differencing, are seen to remove bias most consistently across the full range of α values.

The second column in Figure 1 displays the standard deviations of the estimates. In general these are seen to be broadly comparable between all methods, where methods that do not difference suffer from reduced performance for high α. The dip when estimating α for low values with standard methods is because of boundary effects in the optimization. Here the estimate of α is not permitted to go below 0.5, and the optimization typically converges to the lower bound of 0.5 when $\alpha \leq 0.7$ such that the estimate is biased, but not

variable. Note that as we have averaged over 10,000 replicates, the standard error of the bias estimates can be easily observed by dividing the corresponding standard deviations by $\sqrt{10000} = 100$, and we can see that the bias reductions using the de-biased approach are highly significant.

The third column in Figure 1 displays the root mean square error (RMSE), thus combining information from the first two columns. With standard methods, the observed biases are in general larger than the standard deviations, so the shapes of the RMSE curves generally follow that of the absolute biases, except in the instances discussed earlier for the amplitude parameter. The de-biased methods are now seen to be mostly unbiased such that standard deviation is the main contribution to the RMSE. Overall, reductions in error by an order of magnitude are observed in places, and in general the de-biased methods improve upon the standard method with only a few exceptions.

Finally, we aggregate all information in Figure 1 to produce the average mean, standard deviation, and RMSE for each likelihood estimator, which we present in Table 1. Here we have averaged across all parameter estimates of $\{A, \alpha, c\}$, and over the full range of α considered. To not skew the results in favor of the estimation of any particular parameter, we have averaged the *percentage* (rather than absolute) bias, standard deviation, and RMSE across all the results. We can see that of all the estimators, the de-biased Whittle likelihood using the differenced process performs best in each measure. Overall, of the three procedures—de-biasing, tapering and differencing—de-biasing is the single procedure that yields the greatest overall improvement in parameter estimation.

7.2 Comparison with time domain estimators

Time domain $\mathcal{O}(n \log n)$ approximations to maximum likelihood have recently been made in Anitescu et al. (2012) (see also Dutta and Mondal (2015) and Stein et al. (2013)). These

Table 1: Aggregated results from Figure 1, where we average the percentage bias, standard deviation (s.d.), and root mean square error (RMSE) across all estimates of $\{A, \alpha, c\}$, and over the full range of α considered.

Inference Method	Eqn	Bias	s.d.	RMSE
Standard Whittle (periodogram)	(3.2)	23.69%	10.34%	26.66%
De-Biased Whittle (periodogram)	(4.1)	3.96%	12.97%	13.75%
Standard Whittle (tapered)	(6.1)	18.11%	12.23%	23.12%
De-Biased Whittle (tapered)	(6.2)	2.60%	14.15%	14.41%
Standard Whittle (differenced)	(6.3)	18.99%	9.33%	22.09%
De-Biased Whittle (differenced)	(6.4)	1.19%	8.90%	8.99%

methods use Hutchinson trace estimators (Hutchinson 1990) and circulant embedding, thus removing the need to calculate a matrix inverse or determinant. We contrast our approaches here using the MATLAB package "ScalaGauss" supplied by those authors at `http://press3.mcs.anl.gov/scala-gauss/software/`. We use the parameters selected in their example code for a Matérn process, which only estimates the damping parameter, and assumes the slope parameter is known, and the amplitude parameter is known up to a proportion of the damping parameter. The parameters used, when transformed into the form of equation (7.1), are $A = 1.7725c, \alpha = 1.5, c = 0.0197$ and $n = 1,024$. As the slope parameter is high, we fit the de-biased Whittle likelihood to the differenced process, motivated by the findings of Figure 1. We perform 10,000 repeats and report the results in Table 2. We also include results for maximum likelihood and standard Whittle likelihood, as well as for standard and de-biased tapered likelihoods, where standard methods are not differenced, for comparison.

Table 2: Performance of different inference methods when estimating the damping parameter, α, from a sampled Matérn process. The experiment is repeated over 10,000 independently generated Matérn series of length $n = 1,024$ with parameters, $A = 1.7725c, \alpha = 1.5, c = 0.0197$. The average bias, standard deviation (s.d.), and root mean squared error (RMSE), for each parameter estimate is expressed as a percentage of the true value. Inference methods tested include maximum likelihood, standard/de-biased Whittle likelihood using either the periodogram or dpss taper, and also a normal and fast version of the Anitescu et al. (2012) estimator. CPU times are as performed on a 2.8 GHz Intel Core i7 processor.

Inference Method	Eqn	Bias	s.d.	RMSE	CPU (sec.)
Maximum likelihood	(3.1)	0.029%	2.204%	2.204%	4.257
Standard Whittle (periodogram)	(3.2)	107.735%	101.357%	147.916%	0.139
De-Biased Whittle (differenced)	(6.4)	0.030%	2.212%	2.212%	0.157
Standard Whittle (tapered)	(6.1)	25.550%	20.459%	32.731%	0.168
De-Biased Whittle (tapered)	(6.2)	0.023%	2.558%	2.558%	0.198
Anitescu et al. (normal version)		0.029%	2.205%	2.205%	1.998
Anitescu et al. (faster version)		0.035%	2.223%	2.223%	0.438

The de-biased Whittle likelihood and the method of Anitescu et al. (2012) return estimation errors that are very close to the optimal maximum likelihood. Standard Whittle likelihood performs extremely poorly due to the blurring effects of using the periodogram. The method of Anitescu et al. (2012) requires an order of magnitude more processing time than the de-biased and standard Whittle likelihood. To speed up the Anitescu et al. (2012) method, we have included results with a faster version which uses only 1 Hutchinson trace estimator (as opposed to the 50 used in the example code), but this method is still slower than the de-biased Whittle likelihood and now yields slightly worse estimation accuracy.

Overall, the reported biases are very small compared to the standard deviations, except for standard Whittle likelihood. The standard error of the bias is the standard deviation divided by the square root of the number of replicates, i.e. $\sqrt{10000} = 100$. Therefore the standard errors of the biases are often as large as the biases themselves, and no significance should be placed on the ordering of estimators in terms of bias, other than the poor performance of standard Whittle likelihood. Indeed, this type of Matérn process produces time series that appear nonstationary (as the damping parameter c is extremely small), and as documented by Dahlhaus (1988), standard Whittle likelihood performs poorly in such scenarios. With the de-biased Whittle likelihood however, the properties of the periodogram are corrected, and there is no advantage to be gained by using tapers.

The Anitescu et al. (2012) method is numerically more complicated to implement than the de-biased Whittle likelihood. Whereas our method can be implemented in just a few lines of code, the circulant embedding method requires several more, particularly as random numbers must be generated for the Hutchinson trace estimators. Furthermore, the online code provided in the ScalaGauss package is only for estimating one unknown parameter; it is not clear how the code generalizes to estimate all 3 Matérn parameters, as we have performed in the previous experiment in Figure 1.

8 Conclusions

In this paper we have derived properties of using the Whittle likelihood with finite samples, this yielding unique insights into the governing mechanism of Whittle estimates. To improve the inference procedure, we have proposed the usage of the de-biased Whittle likelihood for estimating the parameters of second-order stationary stochastic processes. The method adjusts the standard Whittle likelihood by replacing the model for the theoretical spectrum with the corresponding "blurred" spectral estimate, which quantifies expected blurring and aliasing artifacts in estimating spectra from observed samples. The proposed method significantly reduces bias and estimation error typically observed with standard Whittle likelihood approaches, while still keeping computations highly efficient, as evidenced in a Monte Carlo study comparing with state of the art alternatives. The method is shown to be consistent, using relatively weak assumptions. Furthermore, The de-biased Whittle likelihood can also be combined with tapering or differencing the data, as is commonly performed in the literature, to further reduce bias and error in parameter estimates.

A key methodological direction of future work is to extend the use of the Whittle likelihood to multivariate processes, as well as to nonstationary time series. While the concept of implementing the blurred theoretical spectrum extends relatively easily, finding a form for its efficient computation is in general non-trivial. Furthermore, to establish consistency with such processes, the proof we provide in the appendix would have to be adjusted accordingly.

Acknowledgments

We wish to thank Dr Jeffrey J. Early for useful discussions, and for providing comments and corrections for this paper.

A Proof of Theorem 1: Consistency of the de-biased Whittle likelihood

We prove Theorem 1 via three propositions. For conciseness, in this document only we drop the X subscript notation from J_X and S_X.

Proposition 1. *Assume that the sequence $\{X_t\}$ is a zero-mean second-order stationary process with spectral density $S(\omega)$. Assume that the spectral density of $\{X_t\}$ is bounded above by the finite value $\|S\|_\infty$. Then the covariance of the Discrete Fourier Transform (DFT) of a sample $\{X_t\}_{t=1}^n$, defined by*

$$J(\omega) = \frac{1}{\sqrt{n}} \sum_{t=1}^{n} X_t \exp(-\mathrm{i}\omega t),$$

is bounded, for some choice of $0 < \omega_p < |\delta_\omega - \omega_p|$ (where $\delta_\omega = \omega_1 - \omega_2 \neq 0$), by

$$|C(\omega_1, \omega_2)| = |\mathrm{Cov}\{J(\omega_1), J(\omega_2)\}| \leq \|S\|_\infty \left\{ \frac{1}{j_p} + \frac{2}{\pi^2(|j_1 - j_2| - j_p)} \left[\log(j_p) + \mathcal{O}(1)\right] \right\},$$

where $j_1 = n\omega_1/2\pi$, $j_2 = n\omega_2/2\pi$ and $j_p = n\omega_p/2\pi$.

Proof. We start by defining the Dirichlet kernel as

$$D_n(\omega) = \sum_{t=1}^{n} \exp(-\mathrm{i}\omega t) = \exp(-\mathrm{i}\omega(n+1)/2) \frac{\sin(n\omega/2)}{\sin(\omega/2)}, \quad |D_n(\omega)| = \left| \frac{\sin(n\omega/2)}{\sin(\omega/2)} \right|. \tag{A.1}$$

We will make use of two properties of the Dirichlet kernel. First, we have from (Thomson et al. 2008, Theorem 15.2) that:

$$|D_n(\omega)| \leq \frac{\pi}{|\omega|}, \quad \text{for} \quad |\omega| \leq \pi. \tag{A.2}$$

Second, we have from (Robinson 1994, Lemma 5) that

$$\int_0^{2\pi j/n} |D_n(\omega)| \, d\omega \sim \frac{2}{\pi} \log(j) + \mathcal{O}(1). \tag{A.3}$$

We now quantify the correlation of the Fourier Transform between frequencies ω_1 and ω_2. This takes the form of

$$\begin{aligned} C(\omega_1,\omega_2) &= \frac{1}{2\pi n}\int_{-\pi}^{\pi}\int_{-\pi}^{\pi}\mathbb{E}\{d\Psi_X(\omega')d\Psi_X^*(\omega'')\}D_n\left(\omega_1-\omega'\right)D_n^*\left(\omega_2-\omega''\right) \\ &= \frac{1}{2\pi n}\int_{-\pi}^{\pi}\int_{-\pi}^{\pi}S\left(\omega'\right)\delta\left(\omega'-\omega''\right)D_n\left(\omega_1-\omega'\right)D_n^*\left(\omega_2-\omega''\right)\ d\omega'\,d\omega'' \\ &= \frac{1}{2\pi n}\int_{-\pi}^{\pi}S\left(\omega'\right)D_n\left(\omega_1-\omega'\right)D_n^*\left(\omega_2-\omega'\right)\ d\omega'. \end{aligned}$$

Without loss of generality we take $\delta_\omega = \omega_1 - \omega_2 > 0$, which constrains $j_1 > j_2$. We assume that $S\left(\omega\right)$ can be upper bounded by $\|S\|_\infty$, i.e. $\|S\|_\infty = \sup_\omega S(\omega)$, and therefore note that

$$|C(\omega_1,\omega_2)| \leq \|S\|_\infty \frac{1}{2\pi n}\int_{-\pi}^{\pi}\left|D_n\left(\omega_1-\omega'\right)D_n^*\left(\omega_2-\omega'\right)\right|\ d\omega'.$$

We now implement a change of variables, recalling the periodicity of the integrand, rewriting this integral as

$$|C(\omega_1,\omega_2)| \leq \|S\|_\infty \frac{1}{2\pi n}\int_{-\pi}^{\pi}\left|D_n\left(\omega''\right)D_n^*\left(\omega''-\delta\omega\right)\right|\ d\omega''.$$

We note that because we have assumed that δ_ω is sufficiently large we can split up the range of integration. Specifically we define the range

$$\begin{aligned} \Omega = (-\pi,\pi) &= (-\pi,-\omega_p)\cup(-\omega_p,\omega_p)\cup(\omega_p,\delta_\omega-\omega_p)\cup(\delta_\omega-\omega_p,\delta_\omega+\omega_p)\cup(\delta_\omega+\omega_p,\pi) \\ &= \cup_{j=1}^{5}\Omega_j. \end{aligned}$$

Note that ω_p is a free variable and only enters into our method of bounding, conditional on the choice of j_1 and j_2. We therefore have that we can split up the integral as

$$\begin{aligned} I &= \frac{1}{2\pi n}\int_{\Omega}\left|D_n\left(\omega''\right)D_n^*\left(\omega''-\delta_\omega\right)\right|\ d\omega'' = \sum_{j=1}^{5}\int_{\Omega_j}\frac{1}{2\pi n}\left|D_n\left(\omega''\right)D_n^*\left(\omega''-\delta_\omega\right)\right|\ d\omega'' \\ &= \sum_{j=1}^{5} I_j. \end{aligned}$$

We shall now bound these individual contributions one-by-one where we use combinations of the relationships given in equations (A.2) and (A.3). We start by considering

$$I_1 = \frac{1}{2\pi n}\int_{-\pi}^{-\omega_p} \left|D_n\left(\omega\right) D_n^*\left(\omega-\delta_\omega\right)\right| \, d\omega = \frac{1}{2\pi n}\int_{\omega_p}^{\pi} \left|D_n\left(\omega'\right) D_n^*\left(\omega'+\delta_\omega\right)\right| \, d\omega'$$
$$\leq \frac{1}{2\pi n}\int_{\omega_p}^{\pi-\delta_\omega} \frac{\pi}{\omega'}\frac{\pi}{\left(\omega'+\delta_\omega\right)} \, d\omega' + \frac{1}{2\pi n}\int_{\pi-\delta_\omega}^{\pi} \left|D_n\left(\omega'\right) D_n^*\left(\omega'+\delta_\omega\right)\right| \, d\omega' = I_{11} + I_{12}.$$

For the first term I_{11} we use condition (A.2)

$$I_{11} = \frac{\pi}{2n}\int_{\omega_p}^{\pi-\delta_\omega} \frac{1}{\omega'}\frac{1}{\left(\omega'+\delta_\omega\right)} \, d\omega' = \frac{\pi}{2n\delta_\omega}\int_{\omega_p}^{\pi-\delta_\omega} \left[\frac{1}{\omega'} - \frac{1}{\left(\omega'+\delta_\omega\right)}\right] \, d\omega'$$
$$= \frac{\pi}{2n\delta_\omega}\left[\log\frac{\pi-\delta_\omega}{\pi} - \log\frac{\omega_p}{\omega_p+\delta_\omega}\right] \leq \frac{\pi}{2n\delta_\omega}\left[\log\left(1+\frac{\delta_\omega}{\omega_p}\right) - \log\frac{\pi}{\pi-\delta_\omega}\right]$$
$$\leq \frac{\pi}{2n\delta_\omega}\left[\frac{\delta_\omega}{\omega_p} - \log\frac{\pi}{\pi-\delta_\omega}\right] = \frac{\pi}{2n\omega_p} - \frac{\pi}{2n\delta_\omega}\log\frac{\pi}{\pi-\delta_\omega} = \frac{1}{4j_p} - \frac{\pi}{2n\delta_\omega}\log\frac{\pi}{\pi-\delta_\omega}.$$

For the next term I_{12} we use the periodicity of the Dirichlet kernel and condition (A.2)

$$I_{12} = \frac{1}{2\pi n}\int_{\pi-\delta_\omega}^{\pi} \left|D_n\left(\omega'\right) D_n^*\left(\omega'+\delta_\omega-2\pi\right)\right| \, d\omega' \leq \frac{1}{2\pi n}\int_{\pi-\delta_\omega}^{\pi} \frac{\pi}{\omega'}\frac{\pi}{\left(2\pi-\omega'-\delta_\omega\right)} \, d\omega'$$
$$= \frac{\pi}{2n}\int_{\pi-\delta_\omega}^{\pi} \frac{1}{\omega'}\frac{1}{\left(2\pi-\omega'-\delta_\omega\right)} \, d\omega' = \frac{\pi}{2n}\frac{1}{2\pi-\delta_\omega}\int_{\pi-\delta_\omega}^{\pi} \left[\frac{1}{\omega'} + \frac{1}{\left(2\pi-\omega'-\delta_\omega\right)}\right] \, d\omega'$$
$$= \frac{\pi}{2n}\frac{1}{2\pi-\delta_\omega}\left[\log\frac{\pi}{\pi-\delta_\omega} - \log\frac{\pi-\delta_\omega}{\pi}\right] = \frac{\pi}{n(2\pi-\delta_\omega)}\log\frac{\pi}{\pi-\delta_\omega}.$$

This establishes the behaviour of the first integral. For I_2 we use both (A.2) and (A.3)

$$I_2 = \frac{1}{2\pi n}\int_{-\omega_p}^{\omega_p} \left|D_n\left(\omega\right) D_n^*\left(\omega-\delta_\omega\right)\right| \, d\omega \leq \frac{1}{2\pi n}\frac{\pi}{\left(\delta_\omega-\omega_p\right)}\int_{-\omega_p}^{\omega_p} \left|D_n\left(\omega\right)\right| \, d\omega$$
$$\leq \frac{1}{4\pi(|j_1-j_2|-j_p)}\left[\frac{4}{\pi}\log(j_p) + \mathcal{O}(1)\right] \leq \frac{1}{\pi^2(|j_1-j_2|-j_p)}\left[\log(j_p) + \mathcal{O}(1)\right].$$

The bound on the third term uses condition (A.2)

$$\begin{aligned} I_3 &= \frac{1}{2\pi n}\int_{\omega_p}^{\delta_\omega-\omega_p} \left|D_n\left(\omega\right)D_n^*\left(\omega-\delta_\omega\right)\right| \, d\omega \leq \frac{1}{2\pi n}\int_{\omega_p}^{\delta_\omega-\omega_p} \frac{\pi}{\omega}\frac{\pi}{\left(\delta_\omega-\omega\right)} \, d\omega \\ &= \frac{\pi}{2n\delta_\omega}\int_{\omega_p}^{\delta_\omega-\omega_p}\left[\frac{1}{\left(\delta_\omega-\omega\right)}+\frac{1}{\omega}\right] \, d\omega = \frac{\pi}{2n\delta_\omega}\left[-\log\left(\frac{\omega_p}{\delta_\omega-\omega_p}\right)+\log\left(\frac{\delta_\omega-\omega_p}{\omega_p}\right)\right] \\ &= \frac{\pi}{n\delta_\omega}\log\left(\frac{\delta_\omega-\omega_p}{\omega_p}\right) = \frac{\pi}{n\delta_\omega}\log\left(\frac{\delta_\omega}{\omega_p}-1\right) \leq \frac{\pi}{n\omega_p} = \frac{1}{2j_p}. \end{aligned}$$

The fourth term resembles the second term, and thus takes the form of

$$\begin{aligned} I_4 &= \frac{1}{2\pi n}\int_{\delta_\omega-\omega_p}^{\delta_\omega+\omega_p} \left|D_n\left(\omega\right)D_n^*\left(\omega-\delta_\omega\right)\right| \, d\omega = \frac{1}{2\pi n}\int_{-\omega_p}^{\omega_p} \left|D_n\left(\omega'\right)D_n^*\left(\omega'+\delta_\omega\right)\right| \, d\omega' \\ &= \frac{1}{2\pi n}\int_{-\omega_p}^{\omega_p} \left|D_n\left(\omega''\right)D_n^*\left(\omega''-\delta_\omega\right)\right| \, d\omega'' = I_2 \leq \frac{1}{\pi^2(|j_1-j_2|-j_p)}\left[\log(j_p)+\mathcal{O}(1)\right]. \end{aligned}$$

Finally, we have that the fifth integral takes the form of

$$\begin{aligned} I_5 &= \frac{1}{2\pi n}\int_{\delta_\omega+\omega_p}^{\pi} \left|D_n\left(\omega\right)D_n^*\left(\omega-\delta_\omega\right)\right| \, d\omega = \frac{1}{2\pi n}\int_{\omega_p}^{\pi-\delta_\omega} \left|D_n\left(\omega'\right)D_n^*\left(\omega'+\delta_\omega\right)\right| \, d\omega' \\ &= I_{11} \leq \frac{1}{4j_p}-\frac{\pi}{2n\delta_\omega}\log\frac{\pi}{\pi-\delta_\omega}. \end{aligned}$$

As $\delta_\omega < \pi$, the $\log\frac{\pi}{\pi-\delta_\omega}$ terms in I_{11}, I_{12} and I_5 sum to a negative value and can hence be ignored. The proposition follows by summing the remaining terms in the integral. □

Proposition 2. *Assume that the sequence $\{X_t\}$ is a zero-mean second-order stationary process with spectral density $S(\omega)$. Assume that the spectral density of X_t is bounded above by the finite value $\|S\|_\infty$. For a sample $\{X_t\}_{t=1}^n$, and a given choice of $\omega_p = 2\pi j_p/n$, linear combinations of the periodogram have a variance that aggregates according to*

$$\mathrm{Var}\left\{\frac{1}{n}\sum_{j=1}^{\lfloor n/2\rfloor} a_j\hat{S}\left(\frac{2\pi j_1}{n}\right)\right\} \leq a_{\max}^2\|S\|_\infty^2\left[\frac{1}{2n}+\frac{2j_p}{n}+\frac{1}{j_p^2}+\frac{4\left[\log^2(j_p)+\mathcal{O}(\log(j_p))\right]}{3n\pi^2}\right].$$

Furthermore the optimal choice corresponds to $j_p = n^{1/3}$ in which case we obtain that

$$\operatorname{Var}\left\{\frac{1}{n}\sum_{j=1}^{\lfloor n/2\rfloor} a_j \hat{S}\left(\frac{2\pi j_1}{n}\right)\right\} \leq a_{\max}^2 \|S\|_\infty^2 \left[\frac{1}{2n} + \frac{3}{n^{2/3}} + \frac{4\left[\log^2(n^{1/3}) + \mathcal{O}(\log(n^{1/3}))\right]}{3n\pi^2}\right]$$
$$= a_{\max}^2 \|S\|_\infty^2 \frac{3}{n^{2/3}} \left\{1 + o(1)\right\}.$$

Proof.

$$\operatorname{Var}\left\{\frac{1}{n}\sum_{j_1=1}^{\lfloor n/2\rfloor} a_{j_1} \hat{S}\left(\frac{2\pi j_1}{n}\right)\right\} = \frac{1}{n^2}\sum_{j_1=1}^{\lfloor n/2\rfloor}\sum_{j_2=1}^{\lfloor n/2\rfloor} a_{j_1} a_{j_2} \operatorname{Cov}\left\{\hat{S}\left(\frac{2\pi j_1}{n}\right), \hat{S}\left(\frac{2\pi j_2}{n}\right)\right\}$$
$$= \frac{1}{n^2}\sum_{j_1=1}^{\lfloor n/2\rfloor} a_{j_1}^2 \operatorname{Var}\left\{\hat{S}\left(\frac{2\pi j_1}{n}\right)\right\} + \frac{2}{n^2}\sum_{j_1>j_2}^{\lfloor n/2\rfloor} a_{j_1} a_{j_2} \operatorname{Cov}\left\{\hat{S}\left(\frac{2\pi j_1}{n}\right), \hat{S}\left(\frac{2\pi j_2}{n}\right)\right\}.$$

We shall now determine what this aggregates to. We start by writing

$$C_{12} = \frac{1}{n^2}\sum_{j_1=1}^{\lfloor n/2\rfloor}\sum_{j_2=j_1+1}^{\lfloor n/2\rfloor} a_{j_1} a_{j_2} \operatorname{Cov}\left\{\hat{S}\left(\frac{2\pi j_1}{n}\right), \hat{S}\left(\frac{2\pi j_2}{n}\right)\right\}$$
$$\leq \frac{a_{\max}^2}{n^2}\sum_{j_1=1}^{\lfloor n/2\rfloor}\sum_{j_2=j_1+1}^{\lfloor n/2\rfloor} \operatorname{Cov}\left\{\hat{S}\left(\frac{2\pi j_1}{n}\right), \hat{S}\left(\frac{2\pi j_2}{n}\right)\right\}$$
$$|C_{12}| \leq \frac{a_{\max}^2}{n^2}\sum_{j_1=1}^{\lfloor n/2\rfloor}\sum_{j_2=j_1+1}^{\lfloor n/2\rfloor} \left|\operatorname{Cov}\left\{\hat{S}\left(\frac{2\pi j_1}{n}\right), \hat{S}\left(\frac{2\pi j_2}{n}\right)\right\}\right|.$$

We now need to bound the covariance of the spectrum, using the covariance of the DFT. We find that

$$\operatorname{Cov}\left\{\hat{S}\left(\frac{2\pi j_1}{n}\right), \hat{S}\left(\frac{2\pi j_2}{n}\right)\right\} = \mathbb{E}\left\{\hat{S}\left(\frac{2\pi j_1}{n}\right)\hat{S}\left(\frac{2\pi j_2}{n}\right)\right\} - \mathbb{E}\left\{\hat{S}\left(\frac{2\pi j_1}{n}\right)\right\}\mathbb{E}\left\{\hat{S}\left(\frac{2\pi j_2}{n}\right)\right\}.$$

Using the fact that the Fourier transform of a Gaussian process is also Gaussian we may use Isserlis' theorem and so obtain that

$$\begin{aligned}\mathbb{E}\left\{\hat{S}\left(\frac{2\pi j_1}{n}\right)\hat{S}\left(\frac{2\pi j_2}{n}\right)\right\} &= \mathbb{E}\left\{J\left(\frac{2\pi j_1}{n}\right)J^*\left(\frac{2\pi j_1}{n}\right)J\left(\frac{2\pi j_2}{n}\right)J^*\left(\frac{2\pi j_2}{n}\right)\right\} \\ &= \mathbb{E}\left\{J\left(\frac{2\pi j_1}{n}\right)J^*\left(\frac{2\pi j_1}{n}\right)\right\}\mathbb{E}\left\{J\left(\frac{2\pi j_2}{n}\right)J^*\left(\frac{2\pi j_2}{n}\right)\right\} \\ &\quad + \mathbb{E}\left\{J\left(\frac{2\pi j_1}{n}\right)J^*\left(\frac{2\pi j_2}{n}\right)\right\}\mathbb{E}\left\{J\left(\frac{2\pi j_2}{n}\right)J^*\left(\frac{2\pi j_1}{n}\right)\right\}. \qquad \text{(A.4)}\end{aligned}$$

Thus it follows that

$$\operatorname{Cov}\left\{\hat{S}\left(\frac{2\pi j_1}{n}\right),\hat{S}\left(\frac{2\pi j_2}{n}\right)\right\} = \left|\operatorname{Cov}\left\{J\left(\frac{2\pi j_1}{n}\right),J\left(\frac{2\pi j_2}{n}\right)\right\}\right|^2. \qquad \text{(A.5)}$$

We therefore find that

$$\begin{aligned}|C_{12}| &\leq \frac{a_{\max}^2}{n^2}\sum_{j_1=1}^{\lfloor n/2\rfloor}\sum_{j_2=j_1+1}^{\lfloor n/2\rfloor}\left|\operatorname{Cov}\left\{J\left(\frac{2\pi j_1}{n}\right),J\left(\frac{2\pi j_2}{n}\right)\right\}\right|^2 \\ &= \frac{a_{\max}^2}{n^2}\sum_{j_1=1}^{\lfloor n/2\rfloor}\sum_{j_2=j_1+1}^{j_1+2j_p}\left|\operatorname{Cov}\left\{J\left(\frac{2\pi j_1}{n}\right),J\left(\frac{2\pi j_2}{n}\right)\right\}\right|^2 \\ &\quad + \frac{a_{\max}^2}{n^2}\sum_{j_1=1}^{\lfloor n/2\rfloor}\sum_{j_2=j_1+2j_p+1}^{\lfloor n/2\rfloor}\left|\operatorname{Cov}\left\{J\left(\frac{2\pi j_1}{n}\right),J\left(\frac{2\pi j_2}{n}\right)\right\}\right|^2 \\ &\leq \frac{a_{\max}^2}{n^2}nj_p\max_{j_1,j_2}\operatorname{Var}\left\{J\left(\frac{2\pi j_1}{n}\right)\right\}\operatorname{Var}\left\{J\left(\frac{2\pi j_2}{n}\right)\right\} + a_{\max}^2\tilde{C}_{12} \\ &\leq \frac{a_{\max}^2}{n}j_p\|S\|_\infty^2 + a_{\max}^2\tilde{C}_{12}.\end{aligned}$$

Using Proposition 1, it follows that

$$\begin{aligned}
\widetilde{C}_{12} &\leq \frac{1}{n^2}\sum_{j_1=1}^{\lfloor n/2\rfloor}\sum_{j_2=j_1+2j_p+1}^{\lfloor n/2\rfloor}\|S\|_\infty^2\left\{\frac{1}{j_p}+\frac{2}{\pi^2(|j_1-j_2|-j_p)}\left[\log(j_p)+\mathcal{O}(1)\right]\right\}^2 \\
&\leq \frac{1}{n^2}\sum_{j_1=1}^{\lfloor n/2\rfloor}\sum_{j_2=j_1+2j_p+1}^{\lfloor n/2\rfloor}\|S\|_\infty^2\left\{\frac{2}{j_p^2}+\frac{8}{\pi^4(|j_1-j_2|-j_p)^2}\left[\log^2(j_p)+\mathcal{O}(\log(j_p))\right]\right\} \\
&= \|S\|_\infty^2\left(\widetilde{C}_{12}^{(1)}+\widetilde{C}_{12}^{(2)}\right).
\end{aligned}$$

First we note that

$$\widetilde{C}_{12}^{(1)} = \frac{1}{n^2}\sum_{j_1=1}^{\lfloor n/2\rfloor}\sum_{j_2=j_1+2j_p+1}^{\lfloor n/2\rfloor}\frac{2}{j_p^2} \leq \frac{1}{2j_p^2}.$$

Then we bound the second term

$$\begin{aligned}
\widetilde{C}_{12}^{(2)} &= \frac{1}{n^2}\sum_{j_1=1}^{\lfloor n/2\rfloor}\sum_{j_2=j_1+2j_p+1}^{\lfloor n/2\rfloor}\frac{8}{\pi^4(|j_1-j_2|-j_p)^2}\left[\log^2(j_p)+\mathcal{O}(\log(j_p))\right] \\
&\leq \frac{8\left[\log^2(j_p)+\mathcal{O}(\log(j_p))\right]}{n^2\pi^4}\sum_{j_1=1}^{\lfloor n/2\rfloor}\sum_{\tau=2j_p+1}^{\lfloor n/2-j_1\rfloor}\frac{1}{(\tau-j_p)^2} \\
&\leq \frac{4\left[\log^2(j_p)+\mathcal{O}(\log(j_p))\right]}{n\pi^4}\sum_{\tau=j_p+1}^{\lfloor n/2-j_p\rfloor}\frac{1}{\tau^2} \leq \frac{4\left[\log^2(j_p)+\mathcal{O}(\log(j_p))\right]}{n\pi^4}\sum_{\tau=1}^{\infty}\frac{1}{\tau^2} \\
&= \frac{4\left[\log^2(j_p)+\mathcal{O}(\log(j_p))\right]}{n\pi^4}\zeta(2) = \frac{4\left[\log^2(j_p)+\mathcal{O}(\log(j_p))\right]}{n\pi^4}\frac{\pi^2}{6},
\end{aligned}$$

where we have used a bound based on the Riemann Zeta function at an argument of 2

$$\zeta(2) = \sum_{j=1}^{\infty} j^{-2} = \frac{\pi^2}{6}.$$

Thus it follows that

$$|C_{12}(j_p)| \leq a_{\max}^2\|S\|_\infty^2\left[\frac{j_p}{n}+\frac{1}{2j_p^2}+\frac{2\left[\log^2(j_p)+\mathcal{O}(\log(j_p))\right]}{3n\pi^2}\right].$$

Finally, as

$$\frac{1}{n^2}\sum_{j_1=1}^{\lfloor n/2\rfloor} a_{j_1}^2 \operatorname{Var}\left\{\hat{S}\left(\frac{2\pi j_1}{n}\right)\right\} \leq a_{\max}^2 \|S\|_\infty^2 \frac{1}{2n},$$

the first statement of the proposition follows. To select the optimal order of j_p we see that:

$$\frac{d}{dx}|C_{12}(x)| = a_{\max}^2\|S\|_\infty^2\left[\frac{1}{n} - \frac{1}{x^3}\right] \Rightarrow x_{\text{opt}}^3 = n, \quad \Rightarrow x_{\text{opt}} = n^{1/3}.$$

The optimal rate is therefore achieved if we select $j_p = \mathcal{O}\left(n^{1/3}\right)$. □

Proposition 3. *Assume that the sequence $\{X_t\}$ is a zero-mean second-order stationary process with spectral density $S(\omega)$. For a sample $\{X_t\}_{t=1}^n$, we define the frequency domain likelihood as*

$$\ell_D(\theta) = -\sum_{\omega\in\Omega}\left[\log\left\{\bar{S}\left(\omega;\theta\right)\right\} + J^H\left(\omega\right)\bar{S}^{-1}\left(\omega;\theta\right)J\left(\omega\right)\right]. \tag{A.6}$$

Assume that we chose $|\Omega| = \Theta(n)$ and that the frequency domain likelihood $\ell_D(\theta)$ has two continuous derivatives in θ. Assume that the spectral density of X_t is bounded above by the finite value $\|S_X\|_\infty$ and below by the non-zero value $S_{\min}$. Then the score $\ell_{f;\theta}\left(\theta\right)$ and the Hessian $\ell_{f;\theta\theta}\left(\theta\right)$, the derivatives of equation (A.6)*, satisfy*

$$\frac{1}{n}\ell_{f;\theta}\left(\theta\right) = \mathcal{O}_P\left(\frac{1}{n^{1/3}}\right), \tag{A.7}$$

and

$$\frac{1}{n}\ell_{f;\theta\theta}\left(\theta\right) + \int_{-\pi}^{\pi}\left(\frac{d\log\bar{S}(\omega;\theta)}{d\theta}\right)^2 d\omega = \mathcal{O}_P\left(\frac{1}{n^{1/3}}\right). \tag{A.8}$$

Proof. We start by noting that

$$\frac{1}{n}\frac{\partial}{\partial\theta}\ell_D(\theta) = -\frac{1}{n}\sum_{\omega\in\Omega}\left[\frac{\partial}{\partial\theta}\log\left\{\bar{S}\left(\omega;\theta\right)\right\} - \frac{\partial\bar{S}\left(\omega;\theta\right)}{\partial\theta}\frac{\hat{S}\left(\omega\right)}{\bar{S}^2\left(\omega;\theta\right)}\right]. \tag{A.9}$$

If we calculate expectations then

$$\mathbb{E}\,\frac{1}{n}\frac{\partial}{\partial\theta}\ell_D(\theta) = -\frac{1}{n}\sum_{\omega\in\Omega}\left[\frac{\partial}{\partial\theta}\log\left\{\bar{S}\left(\omega;\theta\right)\right\} - \frac{\frac{\partial}{\partial\theta}\bar{S}\left(\omega;\theta\right)}{\bar{S}\left(\omega;\theta\right)}\right] = 0. \tag{A.10}$$

Furthermore, using Proposition 2, the variance of the score takes the form of

$$\begin{aligned}\mathrm{Var}\left\{\frac{1}{n}\frac{\partial}{\partial\theta}\ell_D(\theta)\right\} &= \mathrm{Var}\left\{\frac{1}{n}\sum_{\omega}\frac{\partial\bar{S}\left(\omega;\theta\right)}{\partial\theta}\frac{\hat{S}(\omega)}{\bar{S}^2(\omega)}\right\} \\ &\leq \frac{\|S\|_\infty^2\|\frac{\partial\bar{S}}{\partial\theta}\|_\infty^2}{\bar{S}_{\min}^4}\frac{3}{n^{2/3}}\left\{1+o(1)\right\},\end{aligned} \tag{A.11}$$

where $\|\frac{\partial\bar{S}}{\partial\theta}\|_\infty = \max_\omega \frac{\partial\bar{S}(\omega)}{\partial\theta}$ and $\bar{S}_{\min} = \min_\omega \bar{S}(\omega)$. We can therefore conclude using Chebychev's inequality that we can fix $C > 0$ such that

$$\Pr\left\{\frac{1}{n}\left|\frac{\partial}{\partial\theta}\ell_D(\theta)\right| \geq C\frac{\|S\|_\infty\|\frac{\partial\bar{S}}{\partial\theta}\|_\infty}{\bar{S}_{\min}^2}\frac{\sqrt{3}}{n^{1/3}}\left\{1+o(1)\right\}\right\} \leq \frac{1}{C^2}. \tag{A.12}$$

We may therefore deduce that

$$\frac{1}{n}\frac{\partial}{\partial\theta}\ell_D(\theta) = \mathcal{O}_P\left(\frac{1}{n^{1/3}}\right). \tag{A.13}$$

As

$$\begin{aligned}\frac{\partial^2\ell_D(\theta)}{\partial\theta^2} &= -\sum_{\omega\in\Omega}\left[\frac{\partial^2}{\partial\theta^2}\log\left\{\bar{S}\left(\omega;\theta\right)\right\} + J^H\left(\omega\right)\frac{\partial^2}{\partial\theta^2}\bar{S}^{-1}\left(\omega;\theta\right)J\left(\omega\right)\right] \\ &= -\sum_{\omega\in\Omega}\left[-\frac{\left(\frac{\partial\bar{S}(\omega;\theta)}{\partial\theta}\right)^2}{\bar{S}^2\left(\omega;\theta\right)} + \frac{\frac{\partial^2\bar{S}(\omega;\theta)}{\partial\theta^2}}{\bar{S}\left(\omega;\theta\right)}\right] \\ &\quad -\sum_{\omega\in\Omega}\left[2\frac{\left(\frac{\partial\bar{S}(\omega;\theta)}{\partial\theta}\right)^2}{\bar{S}^3\left(\omega;\theta\right)} - \frac{\frac{\partial^2\bar{S}(\omega;\theta)}{\partial\theta^2}}{\bar{S}^2\left(\omega;\theta\right)}\right]\hat{S}\left(\omega\right).\end{aligned} \tag{A.14}$$

Thus we may note that

$$\mathbb{E}\,\frac{1}{n}\frac{\partial^2 \ell_D(\theta)}{\partial \theta^2} = -\frac{1}{n}\sum_{\omega\in\Omega}\frac{\left(\frac{\partial \bar{S}(\omega;\theta)}{\partial \theta}\right)^2}{\bar{S}^2\left(\omega;\theta\right)} \longrightarrow -\int_{\Omega}\frac{\left(\frac{\partial \bar{S}(\omega;\theta)}{\partial \theta}\right)^2}{\bar{S}^2\left(\omega;\theta\right)}\,d\omega. \tag{A.15}$$

Furthermore, we have that,

$$\mathrm{Var}\left\{\frac{1}{n}\frac{\partial^2}{\partial \theta^2}\ell_D(\theta)\right\} = \mathrm{Var}\left\{\sum_{\omega\in\Omega}\left[2\frac{\left(\frac{\partial \bar{S}(\omega;\theta)}{\partial \theta}\right)^2}{\bar{S}^3\left(\omega;\theta\right)} - \frac{\frac{\partial^2 \bar{S}(\omega;\theta)}{\partial \theta^2}}{\bar{S}^2\left(\omega;\theta\right)}\right]\hat{S}\left(\omega\right)\right\}.$$

We define

$$\mathcal{S}\left(\theta\right) = \sup_{\omega}\left|2\frac{\left(\frac{\partial \bar{S}(\omega;\theta)}{\partial \theta}\right)^2}{\bar{S}^3\left(\omega;\theta\right)} - \frac{\frac{\partial^2 \bar{S}(\omega;\theta)}{\partial \theta^2}}{\bar{S}^2\left(\omega;\theta\right)}\right|. \tag{A.16}$$

In this instance we can bound the variance, using Proposition 2, by

$$\mathrm{Var}\left\{\frac{1}{n}\frac{\partial^2}{\partial \theta^2}\ell_D(\theta)\right\} \leq \mathcal{S}^2\left(\theta\right)\|S\|_{\infty}^2\frac{3}{n^{2/3}}\left\{1+o(1)\right\}.$$

We can therefore again conclude using Chebychev's inequality that

$$\Pr\left\{\frac{1}{n}\left|\frac{\partial}{\partial \theta}\ell_D(\theta) + \int_{\Omega}\frac{\left(\frac{\partial \bar{S}(\omega;\theta)}{\partial \theta}\right)^2}{\bar{S}^2\left(\omega;\theta\right)}\,d\omega\right| \geq C\cdot\mathcal{S}\left(\theta\right)\cdot\|S\|_{\infty}\frac{\sqrt{3}}{n^{1/3}}\left\{1+o(1)\right\}\right\} \leq \frac{1}{C^2}. \tag{A.17}$$

This yields the second result. □

We now prove Theorem 1.

Proof. We let $\boldsymbol{\theta}'$ lie in a ball centred at $\boldsymbol{\theta}$ with radius $\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}\|$. We additionally define the Hessian matrix $\mathbf{H}$

$$H_{ij}(\boldsymbol{\theta}) = \frac{\partial^2}{\partial\theta_i \partial\theta_j}\ell_D(\boldsymbol{\theta}). \tag{A.18}$$

Then

$$\frac{1}{n}\nabla\ell_D\left(\hat{\boldsymbol{\theta}}\right) = \frac{1}{n}\nabla\ell_D\left(\boldsymbol{\theta}\right) + \frac{1}{n}\mathbf{H}\left(\boldsymbol{\theta}'\right)\left(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}\right). \tag{A.19}$$

Inverting this equation for $\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}$, and using that $\nabla\ell_D\left(\hat{\boldsymbol{\theta}}\right) = 0$, we obtain with the assumption on continuity on $\frac{\partial^2\ell_D(\boldsymbol{\theta})}{\partial\theta_i\theta_j}$ that

$$\frac{1}{n}\mathbf{H}\left(\boldsymbol{\theta}'\right) = \frac{1}{n}\mathbf{H}\left(\boldsymbol{\theta}\right) + o_P(\mathbf{J}), \tag{A.20}$$

where as usual $\mathbf{J}$ is the matrix of all ones, see (Brockwell and Davis 2009, p.201). Additionally define

$$\mathbb{E}\,\mathbf{H}\left(\boldsymbol{\theta}\right) = \boldsymbol{H}\left(\boldsymbol{\theta}\right). \tag{A.21}$$

Thus it follows that

$$\begin{aligned}
\hat{\boldsymbol{\theta}} - \boldsymbol{\theta} &= -\left[\frac{1}{n}\mathbf{H}\left(\boldsymbol{\theta}\right) + o_P(\mathbf{J})\right]^{-1}\frac{1}{n}\frac{\partial}{\partial\theta}\ell_D\left(\boldsymbol{\theta}\right) \\
&= -\left[\frac{1}{n}\boldsymbol{H}\left(\boldsymbol{\theta}\right) + \mathcal{O}_P\left(\mathbf{J}\frac{1}{n^{1/3}}\right) + o_P(\mathbf{J})\right]^{-1}\mathcal{O}_P\left(\frac{1}{n^{1/3}}\right) \\
&= \left[\frac{1}{n}\boldsymbol{H}\left(\boldsymbol{\theta}\right)\right]^{-1}\left[\mathbf{I} + o_P(1) + \mathcal{O}_P\left(\frac{1}{n^{1/3}}\right)\right]\mathcal{O}_P\left(\frac{1}{n^{1/3}}\right) \\
&= \mathcal{O}_P\left(\frac{1}{n^{1/3}}\right),
\end{aligned} \tag{A.22}$$

using Proposition 3, and applying Slutsky's theorem (Ferguson 1996, p.39), which yields the result. □

References

Anitescu, M., J. Chen, and L. Wang (2012). A matrix-free approach for solving the parametric Gaussian process maximum likelihood problem. *SIAM J. Sci. Comput. 34*(1), A240–A262.

Bloomfield, P. (1976). *Fourier analysis of time series: an introduction*. John Wiley.

Brillinger, D. R. (2001). *Time series: data analysis and theory*. SIAM.

Brockwell, P. J. and R. A. Davis (2009). *Time series: theory and methods*. Springer.

Contreras-Cristan, A., E. Gutiérrez-Peña, and S. G. Walker (2006). A note on Whittle's likelihood. *Commun. Stat.-Simul. C. 35*(4), 857–875.

Dahlhaus, R. (1988). Small sample effects in time series analysis: A new asymptotic theory and a new estimate. *Ann. Stat. 16*(2), 808–841.

Dutta, S. and D. Mondal (2015). An h-likelihood method for spatial mixed linear models based on intrinsic auto-regressions. *J. R. Statist. Soc. B 77*(3), 699–726.

Dzhaparidze, K. O. and A. M. Yaglom (1983). Spectrum parameter estimation in time series analysis. In P. R. Krishnaiah (Ed.), *Developments in Statistics*, pp. 1–96. Academic Press, Inc.

Ferguson, T. S. (1996). *A course in large sample theory*, Volume 49. Chapman & Hall London.

Fuentes, M. (2007). Approximate likelihood for large irregularly spaced spatial data. *J. Am. Stat. Soc. 102*(477), 321–331.

Gneiting, T., W. Kleiber, and M. Schlather (2010). Matérn cross-covariance functions for multivariate random fields. *J. Am. Stat. Soc. 105*(491), 1167–1177.

Granger, C. W. and R. Joyeux (1980). An introduction to long-memory time series models and fractional differencing. *J. Time Ser. Anal. 1*(1), 15–29.

Hutchinson, M. F. (1990). A stochastic estimator of the trace of the influence matrix for Laplacian smoothing splines. *Commun. Stat.-Simul. C. 19*(2), 433–450.

Matsuda, Y. and Y. Yajima (2009). Fourier analysis of irregularly spaced data on R^d. *J. R. Statist. Soc. B 71*(1), 191–217.

Pawitan, Y. and F. O'Sullivan (1994). Nonparametric spectral density estimation using penalized Whittle likelihood. *J. Am. Stat. Soc. 89*(426), 600–610.

Percival, D. B. and A. T. Walden (1993). *Spectral Analysis for Physical Applications: Multitaper and conventional univariate techniques*. Cambridge University Press.

Robinson, P. M. (1994). Rates of convergence and optimal spectral bandwidth for long range dependence. *Probab. Theory Rel. 99*(3), 443–473.

Serroukh, A. and A. T. Walden (2000). Wavelet scale analysis of bivariate time series ii: statistical properties for linear processes. *J. Nonparametr. Stat. 13*(1), 37–56.

Slepian, D. and H. O. Pollak (1961). Prolate spheroidal wave functions, Fourier analysis and uncertainty–I. *Bell Syst. Tech. J. 40*(1), 43–63.

Stein, M. L., J. Chen, and M. Anitescu (2013). Stochastic approximation of score functions for Gaussian processes. *Ann. Appl. Stat. 7*(2), 1162–1191.

Taniguchi, M. (1983). On the second order asymptotic efficiency of estimators of Gaussian ARMA processes. *Ann. Stat. 11*(1), 157–169.

Thomson, B. S., J. B. Bruckner, and A. M. Bruckner (2008). *Elementary Real analysis*. ClassicalRealAnalysis.com.

Thomson, D. J. (1982). Spectrum estimation and harmonic analysis. *Proc. IEEE 70*(9), 1055–1096.

Velasco, C. and P. M. Robinson (2000). Whittle pseudo-maximum likelihood estimation for nonstationary time series. *J. Am. Stat. Soc. 95*(452), 1229–1243.

Wang, T. and Y. Xia (2015). Whittle likelihood estimation of nonlinear autoregressive models with moving average residuals. *J. Am. Stat. Soc. 110*(511), 1083–1099.

Whittle, P. (1953). Estimation and information in stationary time series. *Ark. Mat. 2*(5), 423–434.